

Docket No.: 50277-164

UTILITY PATENT APPLICATION
UNDER 37 CFR 1.53(b)

Box PATENT APPLICATION
Assistant Commissioner for Patents
Washington, DC 20231
Sir:

Transmitted herewith for filing is the patent application of:

INVENTOR: Boris PECHENY
FOR: LEXICAL CACHE

Enclosed are:

- ☒ 27 pages of cover sheet, specification, claims, abstract.
- ☒ Declaration.
- ☐ Priority Claimed.
- ☐ Certified copy of _____
- ☒ 3 sheets of formal drawing.
- ☒ An assignment of the invention to Oracle Corporation
and the assignment recordation fee.
- ☐ An associate power of attorney.
- ☐ A verified statement to establish small entity status under 37 CFR 1.9 and 37 CFR 1.27.
- ☐ Information Disclosure Statement, Form PTO-1449 and reference.
- ☒ Return Receipt Postcard
- ☒ Power of Attorney and Revocation of Previous Powers


The filing fee has been calculated as shown below:

	NO. OF CLAIMS		EXTRA CLAIMS	RATE	AMOUNT
Total Claims	30	-20	10	\$18.00	\$180.00
Independent Claims	4	-3	1	\$78.00	\$78.00
Multiple Dependent Claim(s)					\$0.00
Basic Fee					\$760.00
Total of Above Calculations					\$1,018.00
Less ½ for Small Entity					\$0.00
Assignment & Recording Fee					\$40.00
Total Fee					\$1,058.00

- ☒ Please charge my Deposit Account No. 500417 in the amount of \$1,058.00. A duplicate copy of this sheet is enclosed.
- ☒ The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 500417. A duplicate copy is enclosed.
- ☒ Any additional filing fees required under 37 CFR 1.16.
- ☒ The Commissioner is hereby authorized to charge payment of the following fees during the pendency of this application or credit any overpayment to Deposit Account No. 500417. A duplicate copy of this sheet is enclosed.
- ☒ Any patent application processing fees under 37 CFR 1.17.
- ☒ Any filing fees under 37 CFR 1.16 for presentation of extra claims.

Respectfully submitted,

MCDERMOTT, WILL & EMERY



Stephen C. Carlson
Registration No. 39,929

600 13th Street, N.W.
Washington, DC 20005-3096
(202) 756-8000 SCC:klm
Date: March 8, 1999
Facsimile: (202) 756-8087

RECEIVED 03/08/99

LEXICAL CACHE

FIELD OF THE INVENTION

The present invention relates to computer systems and more particularly to
5 caching techniques for lexical data.

BACKGROUND OF THE INVENTION

Relational databases store information in collections of tables, in which each table
is organized into rows and columns. The various rows of database tables in traditional
applications tend to be accessed with, more or less, a uniform frequency. Thus, the vast
10 majority of accesses to a database table in traditional applications are not skewed to a
relatively small number of rows. Accordingly, various index structures and caches have
been developed for efficiently searching large tables with the assumption that the access
pattern is more or less uniform. Specifically, index structures provide an easily searched
mapping between row identifiers and key values derived from a column of the
15 corresponding row. Many of these index structures, such as a B-tree index, are
characterized by search times that are relatively uniform for each access key.

For applications performing text analysis, on the other hand, the majority of
accesses are highly skewed to relatively few rows of a database table. For example, a
natural language processing application for interpreting English documents may
20 implement a lexicon using a table that contains a row for every English word. The
pattern of accesses to this table is likely to be highly skewed in a Zipf distribution,
because a small percentage of English words (around 10%) account for the vast majority
(>85%) of words in an English document.

Use of conventional relational database index structures to index this table,
25 however, results in a sub-optimal performance for natural language processing
applications, because the search time for very frequently accessed keys is no less than the

access time for rarely accessed keys. What is needed, therefore, is a caching methodology such that searching a table of lexical data for frequently accessed keys results in search times that are significantly smaller than search times for rarely accessed keys. Addressing this need is complicated by the fact that the most frequently used

5 words for a specific topic or set of related documents, aside from a relatively small set of about 1500 words, varies greatly from topic to topic. Therefore, it is difficult to statically determine ahead of time the 40,000-60,000 words to put in a lexical cache for a topic that would have an acceptable hit rate of about 95%.

Furthermore, this need is particularly acute as natural language processing

10 applications grow to access huge tables storing lexical entities such as words and phrases. For example, a lexicon may include 600,000 words and phrases, and the industry trend is toward dramatically larger lexicons.

SUMMARY OF THE INVENTION

Accordingly, these and other needs are addressed by the present invention, which provides a lexical cache comprising a collection of lexical containers, organized according to the length of the words. The present invention stems from the realization that, when the lexicon is divided into several subsets containing words of the same length, the subsets with the shorter length words tend to have a greater the frequency of usage, relative to the number of words in the subset, than subsets with longer length words. For example, words from a subset of 5905 three-letter words are more likely to be used than words from a subset of 10,561 six-letter words.

One aspect of the invention is a computer-implemented method and a computer-readable medium bearing instructions for searching for a string in a lexical cache. In accordance with the methodology, a key is generated based on the string, for example, by compression. A lexical container, such as a hash table, is identified from a plurality of lexical containers based on the length of the key, and the identified lexical container is searched for an entry associated with the string. By identifying the lexical container to be search based on the length of the key, the lexical containers can be implemented easily and efficiently, for example, by a collection of fixed-size key hash tables.

Furthermore, the size and performance of each lexical container can be individually tuned to account for the frequency patterns of each subset of the lexicon divided by length. For example, lexical containers for shorter length words can be configured to be larger than the lexical containers for longer words. Thus, the lexical cache would hold a higher proportion of the subset of shorter length words than the subset for longer length words. Since words from a shorter word length lexical container tend to be more frequently accessed, relative to the size of the lexical container, the lexical cache by its structure will tend to contain more frequently accessed words.

In one embodiment, the string is compressed to generate a key. Based on the length of the key, a hash table is identified from among a plurality of hash tables. The

hash table is organized as sequences of slots for holding respective key values, with each sequence of slots corresponding to a respective hash value. A hash value is computed based on the key, and the hash table is searched based on the hash value for a slot holding a key value matching the key. If a slot having a key value matching the key was found, then the relative position of the key value within the corresponding sequence of slots is moved toward the beginning of the corresponding sequence. By reordering the position of keys in the hash table, more frequently used keys will percolate to the beginning of their sequence, enabling on a dynamic basis faster access times for more frequently used keys.

Still other objects and advantages of the present invention will become readily apparent from the following detailed description, simply by way of illustration of the best mode contemplated of carrying out the invention. As will be realized, the invention is capable of other and different embodiments, and its several details are capable of modifications in various obvious respects, all without departing from the invention.

Accordingly, the drawing and description are to be regarded as illustrative in nature, and not as restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

5 FIG. 1 depicts a computer system that can be used to implement the present invention.

FIG. 2 is a schematic diagram of data structures in accordance with an embodiment of the present invention.

FIG. 3 is a flowchart illustrating how a key is searched for in a lexical cache in
10 accordance with an embodiment of the present invention.

SCANNED

DESCRIPTION OF THE PREFERRED EMBODIMENT

A method and apparatus for caching lexical data is described. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

In a database management system, data is stored in one or more data containers, each container contains records, and the data within each record is organized into one or more fields. In relational database systems, the data containers are referred to as tables, the records are referred to as rows, and the fields are referred to as columns. In object oriented databases, the data containers are referred to as object classes, the records are referred to as objects, and the fields are referred to as attributes. Other database architectures may use other terminology.

Systems that implement the present invention are not limited to any particular type of data container or database architecture. However, for the purpose of explanation, the terminology and examples used herein shall be that typically associated with relational databases. Thus, the terms "table," "row," and "column" shall be used herein to refer respectively to the data container, record, and field.

HARDWARE OVERVIEW

Figure 1 is a block diagram that illustrates a computer system 100 upon which an embodiment of the invention may be implemented. Computer system 100 includes a bus 102 or other communication mechanism for communicating information, and a processor 104 coupled with bus 102 for processing information. Computer system 100 also includes a main memory 106, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 102 for storing information and instructions to be executed

by processor 104. Main memory 106 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 104. Computer system 100 further includes a read only memory (ROM) 108 or other static storage device coupled to bus 102 for storing static information and
5 instructions for processor 104. A storage device 110, such as a magnetic disk or optical disk, is provided and coupled to bus 102 for storing information and instructions.

Computer system 100 may be coupled via bus 102 to a display 112, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 114, including alphanumeric and other keys, is coupled to bus 102 for communicating
10 information and command selections to processor 104. Another type of user input device is cursor control 116, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 104 and for controlling cursor movement on display 112. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the
15 device to specify positions in a plane.

The invention is related to the use of computer system 100 for caching lexical data. According to one embodiment of the invention, caching lexical data is provided by computer system 100 in response to processor 104 executing one or more sequences of one or more instructions contained in main memory 106. Such instructions may be read
20 into main memory 106 from another computer-readable medium, such as storage device 110. Execution of the sequences of instructions contained in main memory 106 causes processor 104 to perform the process steps described herein. One or more processors in a multi-processing arrangement may also be employed to execute the sequences of instructions contained in main memory 106. In alternative embodiments, hard-wired
25 circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 104 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media include, for example, optical or magnetic disks, such as storage device 110. Volatile media include dynamic memory, such as main memory 106. Transmission media include coaxial cables, copper wire and fiber optics, including the wires that comprise bus 102. Transmission media can also take the form of acoustic or light waves, such as those generated during radio frequency (RF) and infrared (IR) data communications. Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, any other magnetic medium, a CD-ROM, DVD, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 104 for execution. For example, the instructions may initially be borne on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 100 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector coupled to bus 102 can receive the data carried in the infrared signal and place the data on bus 102. Bus 102 carries the data to main memory 106, from which processor 104 retrieves and executes the instructions. The instructions received by main memory 106 may optionally be stored on storage device 110 either before or after execution by processor 104.

Computer system 100 also includes a communication interface 118 coupled to bus 102. Communication interface 118 provides a two-way data communication coupling to

a network link 120 that is connected to a local network 122. For example, communication interface 118 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 118 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. 5 Wireless links may also be implemented. In any such implementation, communication interface 118 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 120 typically provides data communication through one or more 10 networks to other data devices. For example, network link 120 may provide a connection through local network 122 to a host computer 124 or to data equipment operated by an Internet Service Provider (ISP) 126. ISP 126 in turn provides data communication services through the worldwide packet data communication network, now commonly referred to as the "Internet" 128. Local network 122 and Internet 128 both use electrical, 15 electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 120 and through communication interface 118, which carry the digital data to and from computer system 100, are exemplary forms of carrier waves transporting the information.

Computer system 100 can send messages and receive data, including program 20 code, through the network(s), network link 120, and communication interface 118. In the Internet example, a server 130 might transmit a requested code for an application program through Internet 128, ISP 126, local network 122 and communication interface 118. In accordance with the invention, one such downloaded application provides for caching lexical data as described herein. The received code may be executed by 25 processor 104 as it is received, and/or stored in storage device 110, or other non-volatile storage for later execution. In this manner, computer system 100 may obtain application code in the form of a carrier wave.

LEXICAL CACHE

Referring to FIG. 2, depicted is a schematic diagram of portions of a lexical cache 200, stored in a computer-readable medium. The lexical cache 200 is an index structure for a lexicon stored in a relational database table (the "lexicon table"). Each row in the
5 lexicon table has a column value that stores a word or phrase.

The lexical cache 200 contains a plurality of entries, such as entry 228, indicating a mapping between (1) row identifiers of the rows in the lexicon table, and (2) key values. For any given row, the key value that maps to the row is derived from the word or phrase stored in the row. Thus, a row in the lexicon table for a particular word is
10 located based on the word by converting to the word into a key and searching the lexical cache 200 for an entry 228 having a key value equal to the search key. If the entry 228 is found, then the row identifier in the entry 228 is used to access the row in the lexicon table. If the entry 228 is not found, then an auxiliary index structure, such as a B-tree built upon the lexicon table, is consulted to determine the appropriate row identifier.
15 Owing to the large size of typical lexicons, it desirable to limit the lexical cache 200 to the most frequently accessed keys value.

The entries of the lexical cache 200 are grouped into units, which are referred to herein as "lexical containers," based on the length of the key. A lexical container is a data structure arranged to store a number of the entries of the lexical cache 200. Stemming
20 from the realization that keys of different lengths have different usage frequencies, each lexical container is dedicated for storing keys of a particular length, except those lexical containers assigned to store entries for a small number of very long but rare keys. In some embodiments, the size of each lexical container will vary depending on the length of the key. For example, it is contemplated that the lexical containers for keys of length 3
25 store more entries in total than lexical containers for keys of length 10.

In a preferred embodiment, the lexical containers are implemented by hash tables 221, such as hash table 220, although it is contemplated that other embodiments may employ other kinds of data structures, such as binary trees or splay trees, to implement the lexical containers. A hash table is a data structure that contains an array of hash table entries. Each hash table entry is associated with a key value. For any given hash table entry, an algorithm is applied to the key value associated with the entry to calculate an index into the array. The index thus produced indicates the location within the array into which the hash table entry should be placed. An array element that holds a key value is called a "slot," and the algorithm that produces the index or "hash value" for a given key is called a "hash function." Various hash functions may be used. According to one embodiment, a hash function is used which applies the byte values of the key as roots to a polynomial and computes a remainder of the sum modulo a predetermined prime number.

The hash function could produce the same hash value for different keys, so that two keys may legitimately be assignable to the same slot. This event is referred to as a "collision." There are a variety of techniques for resolving collisions. One technique, "open address hashing," involves adding a constant to the hash value to index another slot in the hash table. If the new hash value also results in a collision, this process is repeated, generating a sequence of slots, until an empty slot is found, *i.e.* there are no more collisions for the key. "Chaining" is another technique in which linked lists are maintained for each hash value. A linked list maintains, for each entry, a one- or two-byte displacement or other pointer to the next slot in a sequence of slots assigned for the hash value. Upon a collision, the new key value is added somewhere in the linked list.

Preferably, a combination of open address hashing and chaining is used for handling collisions. Specifically, hash table 220 uses open address hashing for the first two slots in sequence of slots for a hash value and chaining for the third and subsequent slots. Accordingly, hash table 220 contains the predetermined prime number of slots in a first region 222 of slots, a second region 224 of slots numbering the predetermined prime

number, and an expansion region 226 of another predetermined number of slots. This combination of open address hashing and chaining guarantees that there are at least two slots in the hash table for every hash value.

The performance characteristics of searching hash table 220 depends on whether the entry for a given key value is cached, and on the number of collisions that are encountered when searching for the entry. The number of collisions for each hash value can be reduced by arbitrarily limiting each sequence of slots allocated to each hash value to a particular maximum. Thus, when the limit is reached, the search among the sequence of slots is terminated as unsuccessful. Empirical testing suggests that the performance of searching a linked list of entries tends to degrade at about 18 entries.

However, if the maximum collision length for each hash value is limited, then it becomes less likely for a given key to be stored in the hash table 220, potentially resulting in an expensive lexical cache 200 miss. Furthermore, increasing the prime number can ameliorate the effects of limiting the collision length, but the prime number cannot be increased beyond the size of the hash table 220. Thus, the desirable performance parameters, the maximum number of collisions and the prime number, depend on the size of the hash table 220.

According to one embodiment, the sizes of the hash tables 221 within lexical cache 200 vary depending on the length of the keys stored in the hash tables 221. Thus, the performance characteristics of hash table 220, characterized by such parameters as the prime number of the hash function, the maximum number of slots, and the maximum number of allowed collisions, are preferably tuned on a case-by-case basis. These and other hash table specific parameters are conveniently stored in an aggregate data structure referred to as a descriptor 216, along with a reference to the corresponding hash table.

The exact values of the tunable parameters for the hash tables will vary from implementation to implementation, depending, for example, on a user-specified lexical cache size in terms of a desired number of total slots. Nevertheless, it is possible to apply

general principles of communications theory in estimating good values for the tunable parameters. See generally, G. K. Zipf, *The Psycho-Biology of Language*, Houghton Mifflin Co., Boston, 1935; C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, 27, 379-423 (1948) and 27, 623-656 (1948).

- 5 For example, Zipf's Law is used to describe the usage frequency distribution among all entries of an English lexicon. Another phenomenon less frequently mentioned or used is the following approximation of the word length – usage frequency dependency:

$$P_{len} = A + KN_{len}e^{-Rlen}, \quad (1)$$

- where P_{len} is the probability of word in a text is of length len , N_{len} is the number of words of length len in the lexicon, and A , K , and R are empirical coefficients. For one English
10 text corpus, the values of the empirical coefficients have been calculated to be: $A=9 \times 10^{-6}$, $K=0.121$, and $R=3.75$.

In one embodiment, therefore, the maximum number of slots, M_i for a hash table storing key values of length i is tuned according to the following formula:

$$15 \quad M_i = \left(\frac{S_a}{N} N_i + \frac{S_b}{R_i} \right) \left(1 + \frac{i}{200} \right), \quad (2)$$

- where $S_a = 0.17S$, $S_b = 0.83S$, S is the desired number of total slots in the lexical cache 200, N is the number of words in a sampled text, N_i is the number of the sampled (compressed) words of length i in the sampled text, R_i is the ratio $\ln(O_i/N_i)/\ln(O/N)$, O is the sum of all occurrences of single words in the sampled text, and O_i is the number of
20 occurrences of single words of length i in the sampled text.

The prime number, P_i for the hash table storing key values of length i is tuned according to the following formula:

$$P_i = \left\lfloor \frac{M_i}{e(2 - R_i/2)} \right\rfloor. \quad (3)$$

- The maximum length of collisions, L_i for the hash table storing key values of
25 length i is tuned according to the following formula:

$$L_i = 2.1 \ln(P_i) - 3.7 R_i. \quad (4)$$

These parameters are calculated for each hash table 221 when the lexical cache 200 is initialized.

In one embodiment, the various lexical containers that the lexical cache 200 comprises are organized into a lexical container matrix 210. Each element of the lexical container matrix 210 references one of the lexical containers. For example, one element 215 in the lexical container matrix 210 contains a pointer to the hash table descriptor 216 of the hash table 220. Hash table descriptor 216, in turn, contains the parameters of hash table 220, and a pointer to hash table 220.

In this embodiment, the element 215 of the lexical container matrix 210 actually references the descriptor 216 for the corresponding hash table 220. The appropriate hash table parameters for the corresponding hash table 220 (for example, the prime number) are fetched from the descriptor 216.

Since each lexical container is associated with keys of a particular length, the lexical container matrix 210 has at least one dimension, which corresponds to the length 212 of the keys stored in the associated lexical containers. Accordingly, the lexical container matrix 210 is indexed in one dimension based on the length of the key. In one embodiment, entries for one-byte keys (not shown) are stored in a 256 element array indexed by the one-byte key, due to the limited number of possible one-byte values. Entries for keys with a length greater than a prescribed cutoff, for example 11, are coalesced into a single row of the lexical container matrix 210. Thus, the lexical containers of the lexical cache 200 are readily identifiable by the length of the key.

The overhead incurred by searching a hash table depends on the number of collisions that are encountered, which is roughly equal to the logarithm of the size of the hash table. In one embodiment, another dimension, “prefix” 214, is added to the lexical container matrix 200 to provide a plurality of different, smaller hash tables for the same key length. Since each hash table is smaller, the number of collisions is fewer and the

performance of the searching one of the smaller tables is improved over searching one large table.

Therefore, each key is assigned to one of the smaller hash tables. In one embodiment, this assignment is based on the “prefix” of the key, which can be defined as a predetermined subset of a particular byte of the key, such as the three least significant bits of the first byte, resulting in a more or less uniform distribution of prefixes.

Accordingly, the prefix 214 dimension of the lexical cache 210 provides a mechanism for identifying one of the smaller hash tables for a particular search key.

SEARCHING THE LEXICAL CACHE

FIG. 3 is a flowchart illustrating how a string representing a word or other lexical item is searched for in the lexical cache 200 in accordance with an embodiment of the present invention. At step 300, the string to be looked up in the lexical cache 200 is compressed to produce the key. A number of compression techniques may be employed, for example, Huffman encoding, n-gram compression, or even no compression. Huffman encoding is a compression technique with variable length codes, and n-gram compression utilizes fixed length codes in which about 80-90 of the possible values of an 8-bit byte are reserved for the basic alphabet and the remaining 170 or so values are assigned to frequently occurring combinations of two or more letters. Compression also helps in obtaining a roughly uniform distribution of prefixes.

At step 302, a descriptor is looked up using the lexical container matrix 210. Specifically, a cell in lexical container matrix is identified based on the length of the key and the prefix (*e.g.* least three significant bits of the first byte of the key). Indexing the lexical container matrix 210 by the length of the key and the prefix yields the cell that contains a reference to a descriptor 216. The descriptor thus referenced identifies the lexical container that would contain the entry for the string for which the search is being performed.

In one embodiment, keys of length one are handled separately in their own 256-entry array (not shown). Key lengths greater than a preset cutoff, e.g. 11, are coalesced into a single row of the lexical container matrix 210 due to the relatively small number of words of such large lengths, especially if compression is employed.

5 At step 304, the hash value for the key is computed based on a hash function specified in the descriptor 216. Specifically, the hash function applies the byte values of the key as roots to a polynomial and computes a remainder of the sum modulo P_i , the prime number stored in the descriptor 216. As a result, the hash value will range from 0 to P_i and will be used to index the first region 222 of slots.

10 At step 306, the hash table is searched until it can be determined if a slot with a key value matching the key is found. If the first slot is empty, then the key is copied into the slot and returned. If the key is equal to the key value of the first slot, then the matching slot is found and returned. If the key is not equal to the key value of the first slot, then the prime number P_i is added to the hash value in an open address fashion to
15 index a slot in the second region 224 of slots.

 Similarly, if the second slot is empty, then the key is copied into the slot and returned. If the key is equal to the key value of the second slot, then the matching slot is found and returned. If the key is not equal to the key value of the second slot, then the links of the current chain is followed in the expansion region until the a slot having a key
20 value matching the key is found or until the maximum length of the chain L_i is reached (e.g. step 310).

 If a matching slot is found in the hash table, the slot is exchanged with the previous slot in the chain (step 308). Thus, upon each access, a slot is moved toward the beginning of the sequence of slots defined by the open address hashing and chaining
25 combination. Consequently, more frequently accessed key values are percolated toward the beginning of the sequence, thereby reducing the number of collisions and improving the search time of future access for such key values.

If the maximum length of the chain L_i is reached, then the key value of the last slot in the chain is replaced by the search key (step 312), thereby discarding a relatively infrequently accessed key value. On the other hand, if the maximum length of the chain L_i is reached, then the search key is simply added as the last slot in the chain (step 314),
5 thereby being cached for a possible future access. Allocating a new slot for the last slot in the chain may be obtained by incrementing a pointer in the descriptor 216 to the next available slot in the expansion region 226 and adjusting the link from the previous slot in the chain.

While this invention has been described in connection with what is presently
10 considered to be the most practical and preferred embodiment, it is to be understood that the invention is not limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

CLAIMS

WHAT IS CLAIMED IS:

- 1 1. A method of searching for a string in a lexical cache, comprising the computer-
2 implemented steps of:
3 generating a key based on the string;
4 identifying a lexical container from among a plurality of lexical containers based on a
5 length of the key; and
6 searching the lexical container for an entry associated with the string.
- 1 2. The method of claim 1, wherein the step of generating a key based on the string
2 includes the step of compressing the string to produce the key.
- 1 3. The method of claim 2, wherein the step of compressing the string to produce the
2 key includes the step of performing an n-gram compression on the string.
- 1 4. The method of claim 1, wherein the step of generating a key based on the string
2 includes the step of using the string as the key.
- 1 5. The method of claim 1, wherein the step of identifying a lexical container includes
2 the steps of:
3 generating a prefix based on the key;
4 identifying the lexical container from among the plurality of the lexical containers
5 based on the length of the key and the prefix.

1 6. The method of claim 1, wherein:
2 the step of identifying a lexical container based on a length of the key includes the
3 step of identifying a hash table based on the length of the key, said hash table
4 containing sequences of slots for holding entries associated with strings, each of
5 said sequences of slots corresponding to a respective hash value; and
6 the step of searching the lexical container for an entry associated with said string
7 includes the steps of:
8 computing a hash value based on the key; and
9 searching the hash table based on the hash value for a slot holding an entry
10 associated with said string.

1 7. The method of claim 6, wherein the step of computing a hash value based on the
2 key includes the step of computing the hash value based on the key and a prime
3 number associated with the hash table.

1 8. The method of claim 7, wherein the step of searching the hash table based on the
2 hash value includes the steps of:
3 indexing one or more fixed regions of the hash table, each of the fixed regions having
4 the prime number of slots, based on the hash value to identify one or more
5 respective slots; and
6 inspecting the one or more respective slots for a respective key value matching the
7 key.

1 9. The method of claim 8, wherein the step of searching the hash table further
2 includes the step of searching for the key in a linked list of slots stored in an expansion
3 region of the hash table, if the key was not found in the one or more respective slots for
4 the key.

1 10. The method of claim 6, further including the step of, if an entry for the string is
2 not found at a first slot that corresponds to the hash value, but is found in a slot that
3 belongs to a sequence of slots that correspond to keys that produce said hash value, then
4 moving a relative position of the entry for the string within the sequence of slots toward
5 the beginning of the sequence of slots.

1 11. The method of claim 6, further comprising the step of initializing a descriptor for
2 the hash table, said descriptor storing a reference to the hash table and parameters for the
3 hash table;
4 wherein the step of identifying a hash table includes the step of identifying a
5 descriptor indicating the hash table and a prime number.

1 12. The method of claim 11, wherein the step of initializing a descriptor for the hash
2 table includes the step of initializing a prime number for use in computing a hash value.

1 13. The method of claim 11, wherein the step of initializing a descriptor for the hash
2 table includes the step of initializing a maximum number of slots for the hash table.

1 14. The method of claim 11, wherein the step of initializing a descriptor for the hash
2 table includes the step of initializing a maximum length of the sequences of slots for the
3 hash table.

15. A method of searching for a string in a lexical cache, comprising the computer-implemented steps of:

compressing the string to generate a key;

identifying a hash table from among a plurality of hash tables based on a length of the

5 key, said hash table containing sequences of slots for holding respective key values, each of said sequences of slots corresponding to a respective hash value;

computing a hash value based on the key;

using said hash value to locate a beginning of the particular sequence of slots that correspond to said hash value;

10 searching the particular sequence of slots for a slot holding a key value matching the key; and

if a slot having a key value matching the key is found in the particular sequence of slots, but is not at the beginning of said particular sequence of slots, then moving a relative position of the key value within the particular sequence of slots toward
15 the beginning of the particular sequence of slots.

1 16. A computer-readable medium bearing instructions for searching for a string in a
2 lexical cache, said instructions arranged, when executed by one or more processors, to
3 cause the one or more processors to perform the steps of:

4 generating a key based on the string;

5 identifying a lexical container from among a plurality of lexical containers based on a
6 length of the key; and

7 searching the lexical container for an entry associated with the string.

1 17. The computer-readable medium of claim 16, wherein the step of generating a key
2 based on the string includes the step of compressing the string to produce the key.

1 18. The computer-readable medium of claim 17, wherein the step of compressing the
2 string to produce the key includes the step of performing an n-gram compression on the
3 string.

1 19. The computer-readable medium of claim 16, wherein the step of generating a key
2 based on the string includes the step of using the string as the key.

1 20. The computer-readable medium of claim 16, wherein the step of identifying a
2 lexical container includes the steps of:
3 generating a prefix based on the key;
4 identifying the lexical container from among the plurality of the lexical containers
5 based on the length of the key and the prefix.

1 21. The computer-readable medium of claim 16, wherein:
2 the step of identifying a lexical container based on a length of the key includes the
3 step of identifying a hash table based on the length of the key, said hash table
4 containing sequences of slots for holding entries associated with strings, each of
5 said sequences of slots corresponding to a respective hash value; and
6 the step of searching the lexical container for an entry associated with said string
7 includes the steps of:
8 computing a hash value based on the key; and
9 searching the hash table based on the hash value for a slot holding an entry
10 associated with said string.

1 22. The computer-readable medium of claim 21, wherein the step of computing a
2 hash value based on the key includes the step of computing the hash value based on the
3 key and a prime number associated with the hash table.

1 23. The computer-readable medium of claim 22, wherein the step of searching the
2 hash table based on the hash value includes the steps of:
3 indexing one or more fixed regions of the hash table, each of the fixed regions having
4 the prime number of slots, based on the hash value to identify one or more
5 respective slots; and
6 inspecting the one or more respective slots for a respective key value matching the
7 key.

1 24. The computer-readable medium of claim 23, wherein the step of searching the
2 hash table further includes the step of searching for the key in a linked list of slots stored
3 in an expansion region of the hash table, if the key was not found in the one or more
4 respective slots for the key.

1 25. The computer-readable medium of claim 21, wherein said instructions are further
2 arranged to cause the one or more processors to perform the step of, if an entry for the
3 string is not found at a first slot that corresponds to the hash value, but is found in a slot
4 that belongs to a sequence of slots that correspond to keys that produce said hash value,
5 then moving a relative position of the entry for the string within the sequence of slots
6 toward the beginning of the sequence of slots.

1 26. The computer-readable medium of claim 21, wherein said instructions are further
2 arranged to cause the one or more processors to perform the step of initializing a
3 descriptor for the hash table, said descriptor storing a reference to the hash table and
4 parameters for the hash table;
5 wherein the step of identifying a hash table includes the step of identifying a
6 descriptor indicating the hash table and a prime number.

1 27. The computer-readable medium of claim 26, wherein the step of initializing a
2 descriptor for the hash table includes the step of initializing a prime number for use in
3 computing a hash value.

1 28. The computer-readable medium of claim 26, wherein the step of initializing a
2 descriptor for the hash table includes the step of initializing a maximum number of slots
3 for the hash table.

1 29. The computer-readable medium of claim 26, wherein the step of initializing a
2 descriptor for the hash table includes the step of initializing a maximum length of the
3 sequences of slots for the hash table.

1 30. A computer-readable medium bearing instructions for searching for a string in a
2 lexical cache, said instructions arranged, when executed by one or more processors, to
3 cause the one or more processors to perform the steps of:
4 compressing the string to generate a key;
5 identifying a hash table from among a plurality of hash tables based on a length of the
6 key, said hash table containing sequences of slots for holding respective key
7 values, each of said sequences of slots corresponding to a respective hash value;
8 computing a hash value based on the key;
9 using said hash value to locate a beginning of the particular sequence of slots that
10 correspond to said hash value;
11 searching the particular sequence of slots for a slot holding a key value matching the
12 key; and
13 if a slot having a key value matching the key is found in the particular sequence of
14 slots, but is not at the beginning of said particular sequence of slots, then moving

- 15 a relative position of the key value within the particular sequence of slots toward
16 the beginning of the particular sequence of slots.

2025 RELEASE UNDER E.O. 14176

LEXICAL CACHE

ABSTRACT OF THE DISCLOSURE

A lexical cache comprises a collection of lexical containers, such as tuned hash table, that are organized according to the length of the keys to be looked up in the lexical
5 cache. In one embodiment, the word is compressed to generate a key. Based on the length of the key and optionally a prefix, a hash table is identified from among the collection of hash tables. A hash value is computed for the key, and the hash table is searched for a slot holding a key value matching the key. If a slot having a key value matching the key was found, then the relative position of the key value within the
10 corresponding sequence of slots is moved toward the beginning of the corresponding sequence.




FIG. 1




FIG. 2

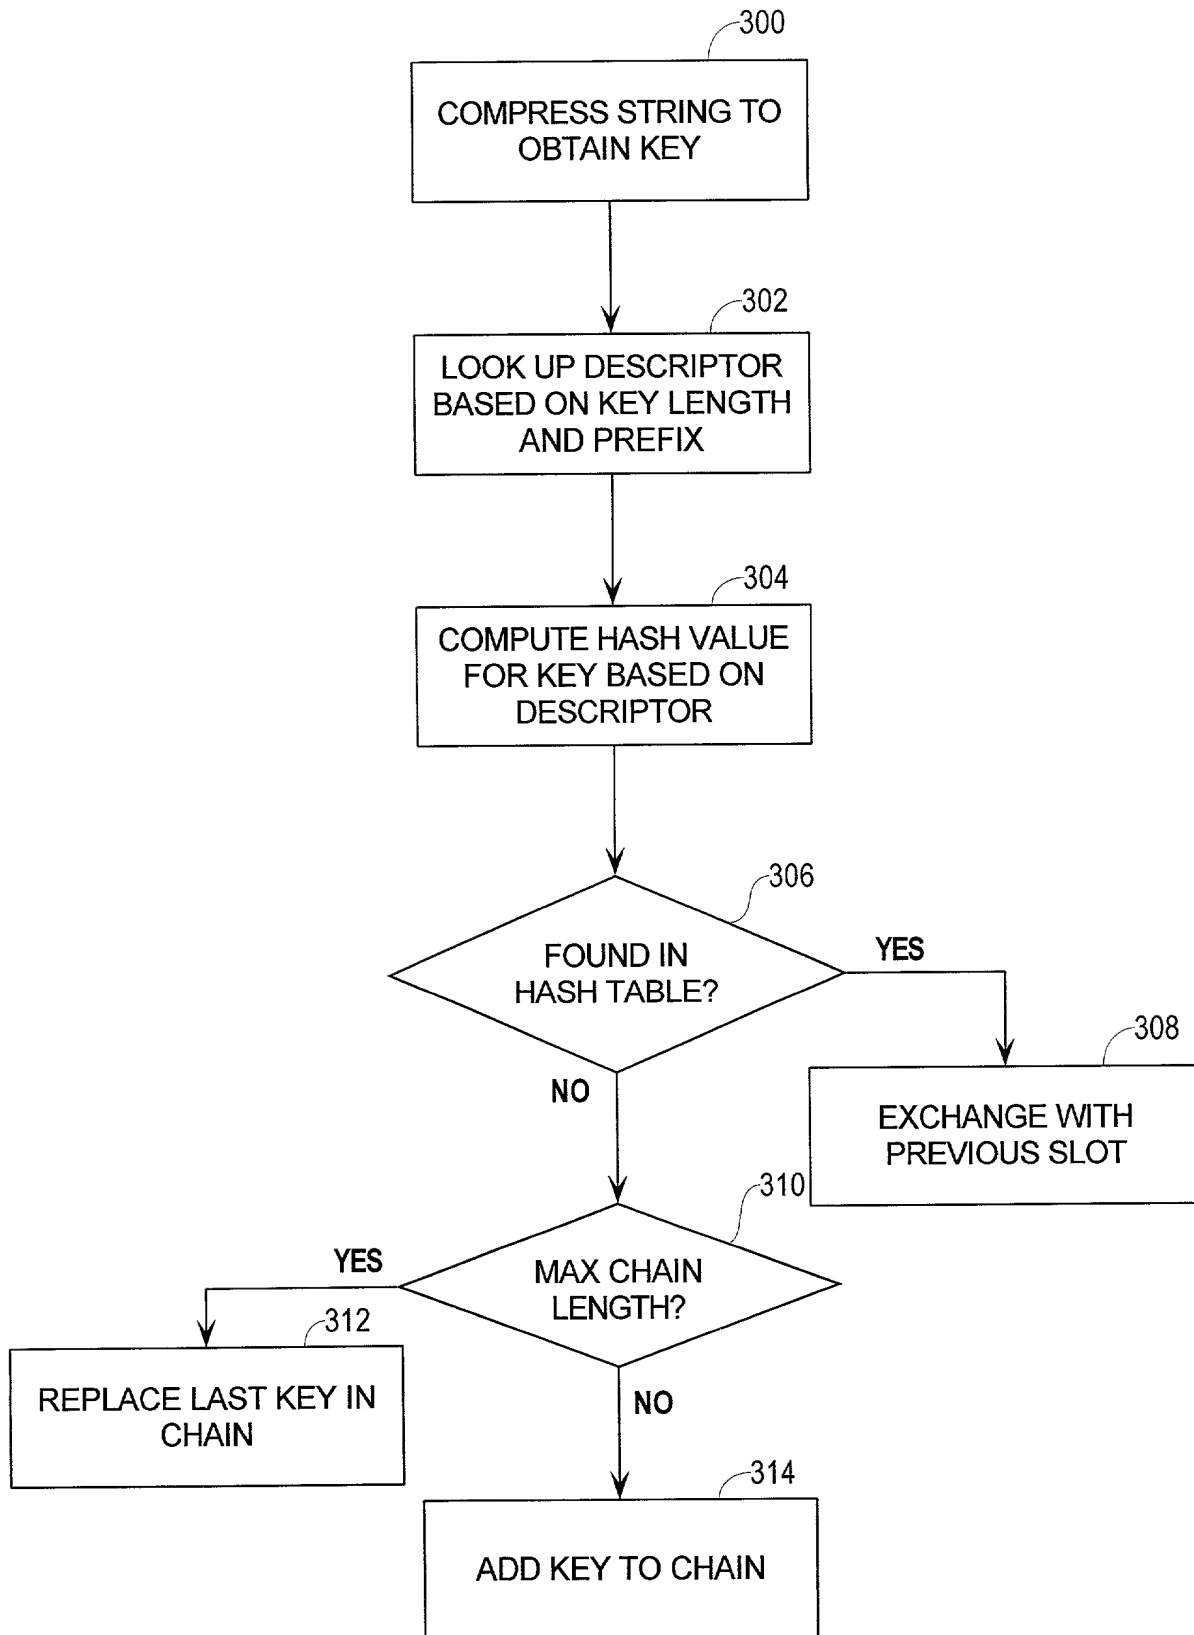


FIG. 3

DECLARATION FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

"LEXICAL CACHE"

the specification of which

 X is attached hereto.
 was filed on _____ as
 United States Application Number _____,
or PCT International Application Number _____,
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims(s), as amended by any amendment referred to above.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 (copy attached).

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), on any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			<u>Priority Claimed</u>	
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<u>Yes</u>	<u>No</u>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<u>Yes</u>	<u>No</u>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<u>Yes</u>	<u>No</u>

I hereby claim the benefit under Title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below

_____ (Application Number)	_____ (Filing Date)
_____ (Application Number)	_____ (Filing Date)

code	name	code	name	code	name	code	name
01	01	01	01	01	01	01	01
02	02	02	02	02	02	02	02
03	03	03	03	03	03	03	03
04	04	04	04	04	04	04	04
05	05	05	05	05	05	05	05
06	06	06	06	06	06	06	06
07	07	07	07	07	07	07	07
08	08	08	08	08	08	08	08
09	09	09	09	09	09	09	09
10	10	10	10	10	10	10	10
11	11	11	11	11	11	11	11
12	12	12	12	12	12	12	12
13	13	13	13	13	13	13	13
14	14	14	14	14	14	14	14
15	15	15	15	15	15	15	15
16	16	16	16	16	16	16	16
17	17	17	17	17	17	17	17
18	18	18	18	18	18	18	18
19	19	19	19	19	19	19	19
20	20	20	20	20	20	20	20
21	21	21	21	21	21	21	21
22	22	22	22	22	22	22	22
23	23	23	23	23	23	23	23
24	24	24	24	24	24	24	24
25	25	25	25	25	25	25	25
26	26	26	26	26	26	26	26
27	27	27	27	27	27	27	27
28	28	28	28	28	28	28	28
29	29	29	29	29	29	29	29
30	30	30	30	30	30	30	30
31	31	31	31	31	31	31	31
32	32	32	32	32	32	32	32
33	33	33	33	33	33	33	33
34	34	34	34	34	34	34	34
35	35	35	35	35	35	35	35
36	36	36	36	36	36	36	36
37	37	37	37	37	37	37	37
38	38	38	38	38	38	38	38
39	39	39	39	39	39	39	39
40	40	40	40	40	40	40	40
41	41	41	41	41	41	41	41
42	42	42	42	42	42	42	42
43	43	43	43	43	43	43	43
44	44	44	44	44	44	44	44
45	45	45	45	45	45	45	45
46	46	46	46	46	46	46	46
47	47	47	47	47	47	47	47
48	48	48	48	48	48	48	48
49	49	49	49	49	49	49	49
50	50	50	50	50	50	50	50
51	51	51	51	51	51	51	51
52	52	52	52	52	52	52	52
53	53	53	53	53	53	53	53
54	54	54	54	54	54	54	54
55	55	55	55	55	55	55	55
56	56	56	56	56	56	56	56
57	57	57	57	57			

(Application Number)	(Filing Date)	(Status - patented, pending, abandoned)
----------------------	---------------	---

Full Name of Sole/First Inventor (given name, family name) BORIS PECHENY

Residence Castro Valley, Calif. Citizenship USA
(City, State) (Country)

2

Title 37, Code of Federal Regulations, Section 1.56
Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

Variable	Mean	Standard deviation	Minimum	Maximum
Age	35.2	12.5	22	65
Gender	0.48	0.50	0	1
Marital status	0.65	0.48	0	1
Education	12.5	2.5	9	16
Income	15.2	8.5	10	25
Health	0.75	0.42	0	1
Smoking	0.25	0.43	0	1
Alcohol	0.15	0.36	0	1
Exercise	0.35	0.48	0	1
Stress	0.60	0.49	0	1
Sleep	0.55	0.50	0	1
Appetite	0.45	0.50	0	1
Mood	0.50	0.50	0	1
Energy	0.40	0.49	0	1
Concentration	0.45	0.50	0	1
Memory	0.40	0.49	0	1
Emotion	0.50	0.50	0	1
Behavior	0.45	0.50	0	1
Thought	0.40	0.49	0	1
Feeling	0.45	0.50	0	1
Perception	0.40	0.49	0	1
Attention	0.45	0.50	0	1
Intuition	0.40	0.49	0	1
Imagination	0.45	0.50	0	1
Reasoning	0.40	0.49	0	1
Logic	0.45	0.50	0	1
Analysis	0.40	0.49	0	1
Synthesis	0.45	0.50	0	1
Evaluation	0.40	0.49	0	1
Comparison	0.45	0.50	0	1
Classification	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem-solving	0.45	0.50	0	1
Decision-making	0.40	0.49	0	1
Communication	0.45	0.50	0	1
Interpersonal	0.40	0.49	0	1
Teamwork	0.45	0.50	0	1
Leadership	0.40	0.49	0	1
Management	0.45	0.50	0	1
Coordination	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem-solving	0.45	0.50	0	1
Decision-making	0.40	0.49	0	1
Communication	0.45	0.50	0	1
Interpersonal	0.40	0.49	0	1
Teamwork	0.45	0.50	0	1
Leadership	0.40	0.49	0	1
Management	0.45	0.50	0	1
Coordination	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem-solving	0.45	0.50	0	1
Decision-making	0.40	0.49	0	1
Communication	0.45	0.50	0	1
Interpersonal	0.40	0.49	0	1
Teamwork	0.45	0.50	0	1
Leadership	0.40	0.49	0	1
Management	0.45	0.50	0	1
Coordination	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem-solving	0.45	0.50	0	1
Decision-making	0.40	0.49	0	1
Communication	0.45	0.50	0	1
Interpersonal	0.40	0.49	0	1
Teamwork	0.45	0.50	0	1
Leadership	0.40	0.49	0	1
Management	0.45	0.50	0	1
Coordination	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem-solving	0.45	0.50	0	1
Decision-making	0.40	0.49	0	1
Communication	0.45	0.50	0	1
Interpersonal	0.40	0.49	0	1
Teamwork	0.45	0.50	0	1
Leadership	0.40	0.49	0	1
Management	0.45	0.50	0	1
Coordination	0.40	0.49	0	1
Organization	0.45	0.50	0	1
Planning	0.40	0.49	0	1
Problem				

4